Network Working Group J Crowcroft (UCL)
Z Wang (UCL)
Internet-Draft A Ghosh (UTS)
<draft-crowcroft-rmfp-00.txt> C Diot (INRIA)
Nov 1996

RMFP: A Reliable Multicast Framing Protocol

Status of this Memo

1. Introduction

There has been considerable interest in reliable multicast, and a number of reliable multicast transport systems have been proposed in the past years.

Reliable multicast transport is considerably more complex than reliable unicast. It is difficult to build a generic reliable transport protocol for muitlcast, much as TCP is a generic transport protocol for unicast, since different applications often have very different reliability requirements and modes of operation.

In this document we propose a framing protocol for reliable multicast transport - Reliable Multicast Framing Protocol (RMFP). RMFP runs over multicast UDP and itself does not provide any reliability (or functionality in a larger extend). Reliability and other protocol functionalities will be defined in specific profiles. The purpose of RMFP is to provides a common framework upon which a set of reliable multicast systems can be built and share similar functionalities where exist.

The philosophy of RMFP is in many respects similar to the one

of RTP. However, RMFP is different from RTP, as we believe that using RTP for reliable multicast is not a right approach and will not lead to a clean application design.

This draft is intended to stimulate more discussion on the one issue of a generic framing protocol for reliable multicast.

2. RMFP Packet Format

RMFP packet header inclcudes common per-packet related fields. An application may include application-specific fields in a preamble header.

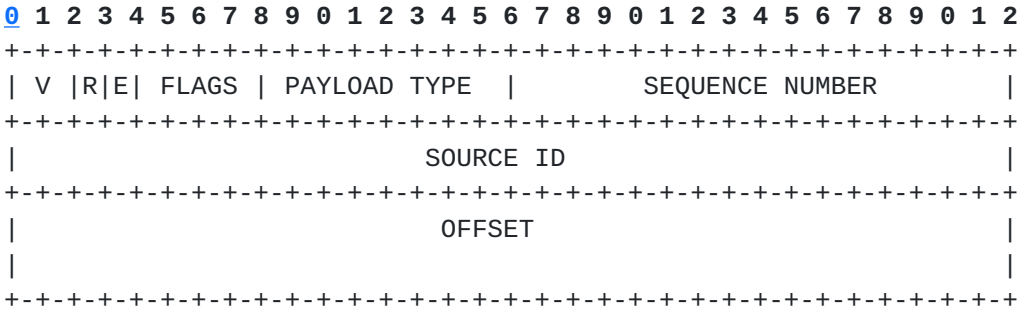

Version(V): 2 bits
    This field identifies the version of RMFP.

Retransmission (R): 1 bit
    This bit, when set, indicates that it is a retransmitted information.

Forward Error Correction (E): 1 bit
    This bit, when set, indicates that FEC is used. The exact format of FEC is determined by Payload Type and its profile.

Flags: 4 bits
    The flags are used for indicating significant features such as object (or ADUs) boundaries. Object boundaries can be used for multiplexing multiple objects within a single session. For example, one can multicast several files within one session.

    0000: reserved
    0001: start mark - the start of an object
    0010: end mark - the end of an object
    other: reserved

Payload Type: 8 bits
    This field identifies the format of the payload and determines its intepretation by the application.
    Profiles will be defined for each payload type.

Sequence Number: 16 bits

The sequence number increments by one for each data packet sent. Sequence number can be used to determine packet losses (including both data packet and retransmitted packets)

Source ID: 32 bits
This field identifies the source. It is generated randomly similar to the SSRC field in RTP. It can be used to detect packet losses.

Offset: 64 bits
This field identifies the position of the data relative to the beginning of the session.

3. RMFP Control Packet Format

RMFP control packets include sender's report packets and receiver's report packets.

Sender's Report Packet

Sender's report is sent periodically by the sender about the data transmitted in the session.

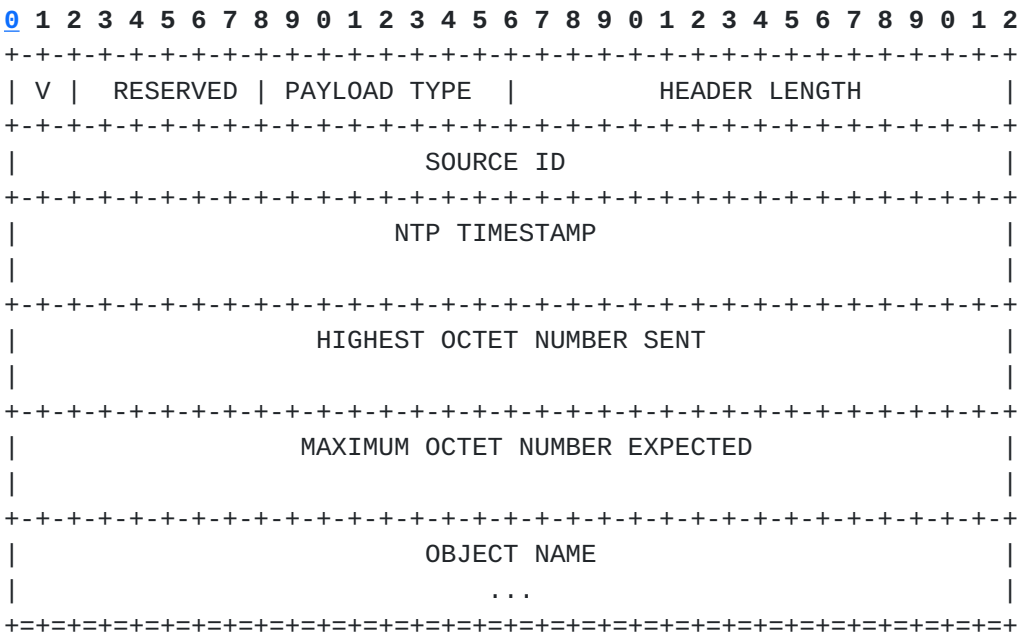

Version(V): 2 bits
This field identifies the version number.

Payload Type: 8 bits
This field is set to xxx for Sender's Report Packets

Header Length: 16 bits
This field specifies the length of the header.

Source ID: 32 bits
    This field identifies the source of the sender

NTP Timestamp: 64 bits
    The NTP timestamp when the report is sent.

Highest Octet Number Sent: 64 bits
    This field indicates the data sent at the time the report
    is sent.

Maximum Octet Number Expected: 64 bits
    This field indicates the total size of the object. An
    application may use the information to allocate space for the
    session. Set to zero if the size is unknown.

Object Name
    This is a variable length field identifying the name of the
    object. It may be a filename, a URL, a message name etc.

Receiver's Report Packet

Receiver's report is periodically sent by the receivers to
give feedback on congestion and packet losses.

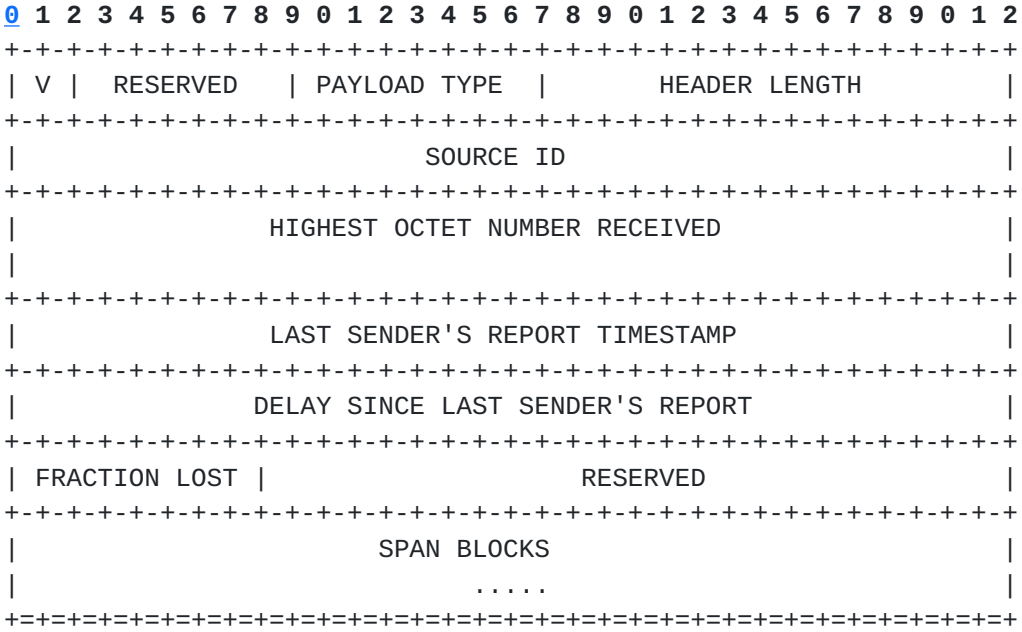

Version(V): 2 bits
    This field identifies the version number.

Payload Type: 8 bits
    This field is set to xxx for Receiver's Report Packets

Header Length: 16 bits
    This field specifies the length of the header.

Source ID: 32 bits
    This field identifies the source of the report

Highest Octet Number Received: 64 bits
    This field indicates the highest octet of the data received so far.

Last Sender's Report Timestamp: 32 bits
    The middle 32 bits of the NTP Timestamp of the most recent
    Sender's Report

Delay Since Last Sender's Report: 32 bits
    The delay, expressed in units of 1/65536 seconds, between
    receiving last Sender's report and sending of this report

Fraction Lost: 8 bits
    The fraction of packets lost since last Sender's report,
    expressed as a fixed point number with the binary
    point at the left edge of the field. Fraction lost is the
    loss rate seen by the receiver. The information may be
    used for congestion control, error recovery purpose by the
    sender.

SPAN Blocks: 64 bits + 32 bits each block
    Each block specifies the offset number and the length of
    a missing data block. The information is used for
    retransmission of lost packets.

4. Open Issues

Profiles for applications

Various and numerous mechanisms can be used to control reliability. Consequently, control information specific to each mechanism cannot be provided ion the RMFP protocol. We propose a profile to be defined for each mechanism. The SRM profile could be based on Parnes' work on reliable RTP, for example. Other profiles could be defined for the various FEC types.

Explicit join/leave

Some reliable applications may need an explicit Join and Leave mechanism. It is not clear to us today how this facility should be provided, or if it has to be provided in RMFP (using reports or a new packet type).

5. Authors's Addresses

J Crowcroft, Zheng Wang
{j.crowcroft, z.wang}@cs.ucl.ac.uk
Department of Computer Science
University College London
Gower Street
London
WC1E 6BT

Atanu Ghosh
atanu@socs.uts.EDU.AU
School of Computing Sciences
University of Technology
Sydney
PO Box 123 , Broadway
NSW 2007
Australia

Christophe Diot
Christophe.Diot@sophia.inria.fr
INRIA
Sophia Antipolis, 2004
route des Lucioles
BP93 06902
France